Acting for the Right Reasons, Abilities, and Obligation[*] [†]

Errol Lord
Franklin & Marshall College
forthcoming in *Oxford Studies in Metaethics, vol. 10*
Winner of the Marc Sanders Prize in Metaethics

1 Introduction

Some moral theorists—we'll call them *objectivists about obligation*—hold that what an agent ought to do is always a function of *all* of the normatively relevant facts.[1] Other moral theorists—we'll call them *perspectivalists about obligation*—hold that what an agent ought to do is rather a function of some *perspectival facts*.[2] Perspectival facts are facts within the agent's perspective. What counts as the agent's perspective is a matter of intermural debate amongst perspectivalists, as is the debate about which facts matter within one's perspective. We'll get to this in due time.

The debate between objectivists and perspectivalists is important because the two views can come apart dramatically. Consider Sick Mother:

[*]Thanks to Kurt Sylvan, Andrew Sepielli, Michael Smith, Tom Kelly, Gideon Rosen, Nat Tabris, Daniel Wodak, Eden Lin. Thanks also to audiences at the Rocky Mountain Ethics Congress, University of Edinburgh, Princeton University and the Georgetown Conference on Reasons and Reasoning, especially Robert Audi, Steve Sverdlik, Liz Harman, Joe Rachiele, Michael Titelbaum, and Dustin Locke.

[†]Draft of September 10, 2013. Citing and quoting encouraged, but please ask permission first (errol.lord@gmail.com).

[1]Theorists who explicitly accept objectivism include Moore (1912), Thomson (1986), and Graham (2010). Moreover, the view is assumed by a prominent tradition in normative theory that focuses on objectivist normative theories. A tacit assumption in much of this work is that what we are morally requried to do is tightly connected to what we ought to do.

[2]A note about how I'm using the word 'obligation.' I am using the word such that *A* is obligated to ϕ *just in case A* ought to ϕ. There is a usage of the word that is somewhat popular amongst moral philosophers where this is not true. Obligations, on this usage, are always things that we *owe* to other agents (see, e.g., Wallace (2012)). I am not using the word in this way. I should say that I mostly use the word 'obligation' to make the prose more elegant. One could replace the word with constructions that just use 'ought,' but that would make this harder to read. Plus, I do think that my use of obligation is also a natural use in English.

Sick Mother

> Jack's mother is in the hospital. She needs an operation in order to survive past this week. Her insurance won't pay. Jack, being a fledgling Art Historian/dealer, doesn't have the money. It looks like his mother is going to die. She will, however, be extremely comforted by Jack's presence in her final days. She lives in California; Jack lives in New York. Jack needs to decide whether to go see her. As it happens, a pawn shop owner is Queens has just unknowingly (and legitimately) bought a rare Picasso. He's selling it at a fraction of the price it's worth. If Jack were to buy it, he would be able to use it as collateral for a loan that would pay for his mother's surgery. The rub, of course, is that he has no idea that this pawn shop even exists, much less that such a deal is to be had there.

Interesting question: What ought Jack do with his day? Ought he go to California or ought he go to Queens? Objectivists hold that all the normatively relevant facts matter when it comes to Jack's obligations. Moreover, it's clear that all the normatively relevant facts point towards going to Queens. He can save his mother that way, and that takes priority one. Perspectivalists, on the other hand, think that only perspectival facts matter when it comes to Jack's obligations. Moreover, it's clear that the perspecitval facts don't support going to Queens; they support going to California. After all, Jack's beliefs, knowledge, and evidence support the thought that it is best to go comfort his mother in her last days. Given his perspective, going to Queens is at best a fool's errand at the cost of his mother dying a lonely death.

This paper is dedicated to arguing for a perspectival view. On this view, the perspectival facts that determine obligation are *possessed normative reasons.* Possessed normative reasons are the normative reasons that are within one's epistemic ken. In §§2-3 I will provide an argument for this view (and thus give an argument against objectivism). The key thought behind the argument is that the facts that obligate must be potentially action-guiding in a certain sense—the facts that obligate must at least potentially be the reasons *for which* we act. This is because when we are obligated to perform some act ϕ, we must at least have the ability to ϕ for the right reasons. The rub will be that we can have the ability to act for the right reasons only if we possess those reasons. This is a huge step forward in a full defense of my view. It also follows that objectivism is false.

Providing the positive argument for my view is not my only goal. I also aim at defusing what I take to be the most compelling objection to perspectivalism. This objection—which goes back to at least Moore (1912) and Ross (1930) and is prominently developed in Thomson (1986) and Graham (2010)—holds that only objectivism can explain the fact that the aim of deliberation is to do what's best in some broad sense.[3] Data often proffered in support of this thought is the fact that it seems like onlookers with more information can have true thoughts about what one ought to do that come apart from what one ought to do given one's perspective.

[3] This is not to assume consequentialism. By 'what's best' I just mean the thing to do given all the facts.

I think my view is compatible with the claim that deliberation aims at what's best. The key thought is that deliberation can aim at what's best even though obligation is *constrained* by our abilities. This is very plausible when it comes to our physiological abilities. Deliberation can aim at what's best even if we aren't every obligated to do what we are physiologically unable to do. I will argue in §4 that we should think of my view in a similar way. My view just enforces a different ability condition. Just like our obligations are constrained by physiological abilities, they are also constrained by some of our agential abilities. This doesn't threaten the claim that the aim of deliberation is to do what's best. Thus, considering the nature of ability constraints more broadly shows that the perspectival view I defend avoids the objection that has long plagued perspectival views.

2 Deliberation, Obligation, Subjective, Objective

2.1 Some Preliminaries

There are two important preliminaries. First, we need to get clearer about the type of obligation that is at issue. We are interested in what I'll call *deliberative obligations*. These obligations are so-called because of their connection to the central deliberative question—viz. what ought I do? For each time the central deliberative questions applies—i.e. every time there is something to be done—there is a correct answer about what is to be done. Of course, this is not to say that there will always be a *single* act that is the act to be done. Often times many actions are permissible. In these cases the correct answer to the central deliberative question will be a disjunction.

The correct answer to the central deliberative question will be the act that you are deliberatively obligated to perform. When there is more than one permissible option, the obligation will scope over the disjunction. The question central to this paper is what determines one's deliberative obligations.

This leads to the second preliminary. It will be helpful to adopt an ideology in order to ask our question in more concrete terms. I will adopt the ideology of *normative reasons*. Normative reasons are facts that count in favor of actions and attitudes. We can frame the debate by appealing to normative reasons. On this framing, objectivists hold that deliberative obligations are a function of all of the normative reasons. Whatever is best supported by all of the reasons is what one ought to do. Perspectivalists hold that only the reasons within one's perspective can determine what you're deliberatively obligated to do. Whatever is best supported by the reasons within your perspective is what you ought to do.

On my perspectival view, (at least part of) your perspective is made up of the normative reasons you *possess*. The normative reasons you possess are the normative reasons that are in your epistemic ken. So for example, while it is true that there is a reason for me go to the store if we're out of milk, I don't possess that reason to go to the store unless I'm aware of the fact that we're out of milk. If the last of the milk is currently being consumed while I'm at the office

(and thus know nothing of our milk situation), then even though there is a reason for me to go to the store, I do not possess that reason to go to the store.

It is, as one might expect, controversial what constitutes one's epistemic ken and thus it is controversial which epistemic relation is constitutive of the possession relation. I would like to be neutral on this here. I think most views of possession are compatible with the views defended in this paper. For simplicity I will assume that the possession relation is knowledge.[4] This is because all of the going views in the literature hold that knowing some reason r is sufficient for possessing r.

2.2 A Dialectical Primer

The dispute between objectivists and perspectivalists has been dominated by two types of cases, which I'll call *simple ignorance cases* and *sophisticated ignorance cases*. In order to know exactly what is at dispute between the two camps, it is helpful to begin by thinking about the dialectic about the cases.

Let's start with the simple cases. We've already seen a simple case. This is Sick Mother. In Sick Mother, Jack is ignorant of the pawn shop and the Picasso. Given his perspective, he ought to go to California. But given all the facts, he ought to go to Queens.

The unifying theme of the simple cases is that the characters in the cases are ignorant of some normatively relevant facts. Given what the characters know, some act ϕ-ing is obligatory. But given all the facts, another act, ψ-ing, is obligatory. I think it is fair to say that most people's initial intuitions about the simple cases support perspectivalism. However, things are more complicated than they might initially seem because the objectivist has a compelling response.

The response has two important parts. The first is that in the simple cases it is always reasonable to believe that the action that is obligatory in light of one's perspective is also permitted by the balance of all the reasons. This is true in Sick Mother. It is reasonable for Jack to think that the balance of all the reasons permits him to go to California. In fact, it's reasonable for him to believe that they require him to do so.

The second part consists in the claim that we should divorce the deontic from the hypological. That is, we shouldn't hold that there are any necessary connections between claims about what ought to be done and claims about what we're praiseworthy or blameworthy for doing or not doing. In particular, claims the objectivist, we shouldn't think that doing wrong is sufficient for being blameworthy—there can be blameless wrongdoing. Importantly for our purposes, there is blameless wrongdoing when one falsely believes that the balance of all the reasons supports ϕ-ing and one ϕ-s. Moreover, we can nicely explain why this is. The ϕ-ing is blameless because it was reasonable to believe that ϕ-ing was supported by the balance of all the reasons. But it was wrong because this belief is *false*.

[4] At the end of the day, I think that being in a position to know is the relevant relation. See (Lord, 2013, ch. 3) for a defense of this view. Again, I don't think it's essential to learning the lessons of this paper that one hold any particular view about possession.

This response is dialectically compelling. This is because it is anchored in the very plausible claim that there can be blameless wrongdoing. Moreover, given objectivism, the simple cases are paradigm cases of blameless wrongdoing. At the very least, this response should dampen the strength of one's intuitions about the simple cases.

Fortunately for perspectivalists, there are the sophisticated cases. The standard objectivist response to the simple cases is not available when it comes to the sophisticated cases. The most famous sophisticated case is Mine Shaft.[5]

Mine Shaft

> A group of 10 miners are trapped in a mine. They are either trapped in shaft *A* or in shaft *B*. It is not known which shaft they are in. Flood waters are approaching the shafts. Billy has the choice to sandbag shaft *A*, sandbag shaft *B*, or not sandbag either. She knows that if she sandbags *A* and the miners are in *A*, all the miners will survive. She knows the same is true of *B*. She also knows that if she sandbags either shaft and the miners are in the other shaft, they will all die. Finally, she knows that if she does nothing, then 9 of the 10 will survive.

First things first, it is very plausible that Billy ought to do nothing. That is, it is very plausible that she ought to do the thing that guarantees that 9 miners will survive. It is simply too risky to sandbag either shaft. At best she will save one life and at worse she will let ten die.[6]

The most important feature of the sophisticated cases when it comes to the current dialectic is that in the sophisticated cases one is not in a position to reasonably believe that the balance of reasons supports the act that is best supported by the facts in one's perspective. In Mine Shaft, Billy *knows* that doing nothing is not the act that will bring about the best outcome. Nevertheless, it seems like she should do nothing.

This blocks the objectivist response to the simple cases from applying to the sophisticated cases. This is because it was crucial to that response that the characters reasonably believe that the balance of all the reasons supports the act that the facts in their perspective supports. The characters in sophisticated cases can't reasonably believe this. Thus, we can't explain why they are blameless in doing the second best option by appealing to a reasonable but false belief. The best explanation of why they are blameless for doing the second best option, it seems, is that they ought to.

2.3 Objective and Subjective

Before moving on, it is important to say something about a common reaction to the above dialectic. The first reaction of many is that we can explain all that needs to be explained by

[5] The case was made famous by Parfit (2011). It orginated in Regan (1980). It is structurally analogous to the doctor case in Jackson (1991) and the envelope case in Ross (2006).

[6] Note that even if you are uncomfortable with the idea of letting anyone die, there are other cases with the same structure that lack this feature (see the doctor case in Jackson (1991) or the envelope case in Ross (2006)).

appealing to the distinction between the objective and subjective obligation. Objective obligations are a function of all the normatively relevant facts. Subjective obligations are in some way perspectival.

The common thought is that in both the simple and sophisticated cases there is something we are objectively obligated to do and something we are subjectively obligated to do. Our intuitions in favor of perspectivalism are really just tracking subjective obligations and our intuitions in favor of objectivism are really just tracking objective obligations. However, neither type of obligation takes precedent. They are just associated with different things of interest to normative theory.[7]

It is important to stress that the appeal to the objective/subjective distinction I am interested in right now is deflationary when it comes to what I've called the ought of deliberation. On the view under consideration, there is *no* conceptual room for the ought of deliberation. There are just the subjective obligations and the objective obligations. Those who hold this view think that the present paper and the dialectic in which it is a part is a mistake. There is just no interesting question to ask about the ought of deliberation. Let's call this view the deflationary view.

The deflationary view has a nice conciliatory tone to it, but I think that it is hard to maintain without committing oneself to some unsavory conclusions. In short, in order to hold this view, one has to think that there are often deontic dilemmas. This is because the objective and subjective obligations will often come apart. All of the ignorance cases are like this. In those cases, defenders of the deflationary view must hold that there are simply two incompatible things one ought to do.

Defenders of the deflationary view will quickly point out that they think the two forms of obligation are *incommensurable* in a certain sense. This is supposed to help with the dilemma objection. And it does to a certain extent. It would be worse if you were obligated in the same sense to perform incompatible actions. However, I don't think it fully dampens the force of the objection. To see this, it's helpful to reflect on sophisticated ignorance cases like Mine Shaft.

Billy knows that doing nothing is the second best option. If she is conceptually sophisticated enough, then she is in a position to know that she objectively ought sandbag one of the shafts. Moreover, if she is conceptually sophisticated enough, she is in a position to know that she subjectively ought to do nothing. If the deflationist is right, these exhaust the deontic facts. But it is very implausible that Billy's has deliberated about all that can be deliberated about. Billy can ask a further question: Which obligation ought I satisfy? But the 'ought' in this question is just the ought of deliberation. This is good reason to think the deflationist is wrong.[8]

Of course, this is not the deny the importance of objective obligation or subjective obligation. Those notions might have important theoretical roles to play. Moreover, it is not to

[7] For a nice presentation of this view, see Schroeder (2009).

[8] For similar remarks, see Jackson (1991); Kolodny & MacFarlane (2010); Zimmerman (2008); Kiesewetter (2011); Graham (2010).

deny the potential relevance of objective or subjective obligations to the debate over deliberative obligations. After all, objectivists hold that the deliberative obligations just are the objective obligations. Moreover, a serious contender holds that the deliberative obligations just are the subjective obligations.[9] If the deliberative obligations reduced to the objective or subjective obligations, then in a sense—a metaphysical one—there would only be the objective and subjective obligations. But in another sense—a conceptual one—there are still three notions. Namely, the objective notion, the subjective notion, and the deliberative notion.

Putting the deflationary view to one side, the rest of the paper will be dedicated to defending my perspectival view.

3 Abilities, Acting for the Right Reasons, and Obligation

It is close to uncontroversial that our abilities can constrain our obligations. Perhaps the least controversial ability condition is the Physiological Abilities Condition:

> **Physiological Ability Condition:** If A ought to ϕ, then A has the physiological ability to ϕ.

The Physiological Ability Condition is supported by strong intuitions. To see this, consider Dunk for Money:

> **Dunk for Money**
>
> Mark Cuban, being the eccentric billionaire that he is, decides to have a raffle. The person whose name is chosen gets a shot at a $10 million prize. In order to win the prize, one has to do a 360° dunk on an NBA regulation sized hoop. Just for fun, Sam enters the raffle. She is the lucky winner. Unfortunately for Sam, she is only 4'11" tall. Because of this, she lacks the ability to dunk on an NBA regulation sized hoop.

It is very plausible that it is not the case that Sam ought to perform the 360° dunk. It's plausible that she ought to *try* to dunk, but not that she ought to dunk. A plausible hypothesis about why this is true is that she lacks the physiological ability to dunk. Sam's physiological abilities seem to be constraining her obligations. If she had the ability to dunk, it would be the case that she ought to. This is because dunking would be best.

The Physiological Abilities Condition is fairly weak. It is compatible with objectivism. I will argue for an abilities condition that is not compatible with objectivism. This is the Right Reasons Abilities Condition:

[9] Perhaps my view is a view like this. In order to find out, we'd have to investigate the essential properties of subjective obligations. I am not interested in doing this here.

> **Right Reasons Ability Condition:** If A ought to ϕ, then A has the ability to ϕ for the right reasons.

The right reasons are the reasons that determine one's obligation. The intuitive idea behind the Right Reasons Ability Condition is that the reasons that determine our obligations can be action guiding. They can be the reasons *for which* we act.

It will be important going forward to have an intuitive grip on the notion of acting for normative reasons. First, let's consider some cases.

Consider Jenny. Jenny needs to get to work. Unfortunately there was a big snow recently. Jenny's car is thus covered in snow. She deliberates about alternative ways of getting to work, but decides—rightly—that she needs to drive. She thus digs her car out. The reason for which she digs her car out is that she needs to get to work. That seems like a good reason to dig her car out.

Now consider Bill. Bill has high cholesterol. He is at a restaurant and is deciding between the fish and the steak. While deliberating about his order, he considers how much he'd enjoy the steak, how much he'd enjoy the fish, and the health considerations. He decides—again, rightly—that his health should guide his choice. Thus, he gets the fish. The reason for which he gets the fish is that he has high cholesterol. Moreover, that seems like a strong reason for him to get the fish.

I think we can tease out of these cases some general truths about acting for the right reasons. The first thing to say is that it is very plausible that we can provide a certain kind of *explanation* of why Jenny digs her car out and of why Bill orders the fish by appealing to the right reasons. Bill gets the fish because of his high cholesterol and Jenny digs her car out because she has to get to work. These explanations aren't merely causal. They are also normative in a particular way. They explain why the actions Bill and Jenny performed are *justified.* Jenny is justified because she needs to get to work and Bill is justified because he has high cholesterol. Let's call explanations of this kind *justificatory explanations*. It's plausible, then, that the Explanatory Condition is true:

> **Explanatory Condition:** If A ϕs for a normative reason r, r provides a justificatory explanation of why A ϕs.

An interesting question to ask is what is it that makes it the case that we can provide a justificatory explanation of why Bill and Jenny did what the did. A plausible answer to this question is that we can provide the justificatory explanation because Bill and Jenny are *sensitive* to the right reasons. They are, that is, sensitive to the fact that the facts for which they act are normative reasons to act in that way. It's plausible to suppose that they wouldn't act the ways in which they did if those facts didn't provide them with normative reasons to act in those ways.[10] They

[10] This is *not* intended to be an analysis of sensitivity. I don't think any counterfactual analysis is adequate. I think that we analyze this sensitivity in dispositional terms (and I don't think we can analyze dispositions counterfactually). The sensistivity involved is the disposition not to perform the action if the reasons were defeated. I think is that interesting counterfactuals are (usually) entailed by the fact that one has these dispositions. See (Lord, 2013, ch. 4) and Lord & Sylvan (MS) for more.

in some way track the relevant normative considerations. And this seems like a very important part of acting for the right reasons. This supports the Sensitivity Condition:

> **Sensitivity Condition:** If A ϕs for a normative reason r, A's ϕ-ing is sensitive to the fact that r is a normative reason to ϕ.

With these thoughts in hand, back to the objectivist. The rub for the objectivist is that I think that in order to meet the Right Reasons Ability Condition for some reason r, one must possess r. But in order to possess r, r has to be within one's perspective. Thus, objectivism is false. Here's the argument in a more perspicuous form:

(1) If A ought to ϕ, then A has the ability to ϕ for the right reasons (Right Reasons Ability Condition).

(2) If A has the ability to ϕ for the right reasons, then A possesses the right reasons.

(C) If A ought to ϕ, then A possesses the right reasons.

If this argument is sound, then objectivism is false and some perspectival view is true. The rest of this section will be dedicated to defending (1) and (2).

3.1 In Defense of the Right Reasons Ability Condition

The intuitive idea behind the Right Reasons Ability Condition is that the reasons that determine our obligations must be able to guide our actions. The paradigm way in which reasons guide our actions is by being the reasons for which we perform actions. I think there are at least two arguments that flesh out this intuition.

The first argument turns on a principle tying credit to acting for the right reasons. It is plausible that token actions are creditworthy only if they are performed for the right reasons. Consider some cases.

Good Husband

> Brandon's wife Jen's birthday is tomorow. She badly wants and needs a new hat. He buys her a new hat. Moreover, the reasons for which he buys her the hat is that she wants it and she needs it.

Bad Husband

> Brandon's wife Jen's birthday is tomorow. She badly wants and needs a new hat. He buys her a new hat. However, the reason for which he buys her the hat is that he dislikes her hair and hopes that she will cover it up with the hat.

In both Good Husband and Bad Husband Brandon ought to buy Jen a new hat.[11] Moreover, in both cases the reasons that obligate Brandon are the facts that Jen wants and needs a new hat for her birthday. The difference between the cases is that it is only in Good Husband that Brandon's action is guided by the reasons that obligate him. It is because of this that it is plausible to think that Brandon's token act is creditworthy only in Good Husband. After all, in Bad Husband it is an accident that Brandon performs the action that he is obligated to perform. That is, it is just a coincidence that what Brandon is motivated to do (for selfish reasons) happens to align with what he is obligated to do. This is not so in Good Husband. Brandon performs the action he does precisely because there are decisive reasons to. This kind of non-accidentality seems required for creditworthiness.[12]

This supports Credit:

> **Credit:** *A*'s ϕ-ing is creditworthy if and only if *A* ϕs for reasons that make ϕ-ing permissible.

In cases where one is obligated to ϕ (and hence no other act is permissible), it follows from Credit that a token ϕ-ing is creditworthy only if it is performed for the reasons that obligate one to ϕ.

Now for the rub. If Credit is true and Right Reasons Ability Condition is false, then there will be cases where one ought to ϕ even though it is impossible for one's token ϕ-ing to be creditworthy. Given the connection between creditworthiness and non-accidentality, this means that there will be cases where one ought to ϕ even though there is no non-accidental way for the action one actually performs to be the action one ought to perform. This is a large cost to pay.

The second argument turns on the relationship between acting for the right reasons and correct deliberation. Recall that there is a tight connection between correct deliberation and our deliberative obligations. If we are obligated to ϕ at *t*, then correct deliberation will end in ϕ-ing at *t*.

It is plausible that there is a tight connection between correct deliberation and acting for the right reasons. Suppose that at *t* one is obligated to ϕ by some reason *r*. Thus, correct deliberation would end in ϕ-ing. Furthermore, it is very plausible that correct deliberation would end in ϕ-ing *based* on *r*. That is, if one were to deliberate in a correct way, one would ϕ because of *r*.

It should be stressed that this result doesn't fall directly out of the minimal connection between deliberation and obligation that we used to glom onto deliberative obligation. The minimal connection was just that if you ought to ϕ, then correct answer to the central deliberative question is: ϕ. This establishes nothing about what fully correct deliberation is like.

[11] In order to not beg any questions against either perspectivalist or objectivists, let's assume Brandon knows all of the relevant facts.

[12] A note about creditworthiness. I am using it in the standard sense where what it is to be creditworthy is for there to be reasons to give you positive credit (or praise) for what you did. So one can be creditworthy even if someone else shouldn't give you positive credit (because the reasons there are to do this are outweighed).

Despite this, it is still very plausible that correct deliberation will run through the reasons that obligate. If one is to deny Right Reasons Ability Condition, one will have to maintain that, in some cases, even though the correct answer to the central deliberative question is 'ϕ' and this is true in virtue of the fact that r decisively supports ϕ-ing, it is not possible for one to ϕ because of r. Again, this seems like a big cost to pay. It is bizarre to think that one ought to ϕ because of r but that r cannot play this motivating role in correct deliberation.

Denying the Right Reasons Ability Condition thus has at least two major costs. First, one has to deny that it is always possible to do what one is obligated to do in a way that is creditworthy. Second, one has to deny that it is always possible to correctly deliberate from the reason that obligates to the action that is obligatory.

It is tempting for the objectivist to reply to these arguments by again appealing to the fact that the deontic comes apart from the hypological. Credit is a hypological notion. Thus, since the deontic doesn't necessarily align with the hypological, we shouldn't expect that it's always possible to do what's obligated in a creditworthy way.

In reply, note that this is a much more radical severing of the deontic from the hypological than the one we already granted. Recall that it is quite plausible that blameworthiness doesn't necessarily co-travel with wrongdoing. That is, it's possible to be blameless even though you've done something wrong. By denying the Right Reasons Ability Condition, one makes a much more radical claim. Rather than saying that in some caes it's *possible* for the deontic and hypological to *come apart*, one is saying that in some cases it's *impossible* for the deontic and hypological to *stay together.* No one denies that in cases where it's possible to be blameless even though you do something wrong, it's also possible to be blameworthy even though you've done something wrong. To deny the Right Reasons Ability Condition is to hold that in some cases it's impossible to do the right thing in a way that is creditworthy. Thus, I don't think the thought I'm willing to grant about the deontic and hypological opens the door very wide for this type of response. Denying the Right Reasons Ability Condition is costlier than that.

3.2 In Defense of (2)

I suspect that objectivists are so far unperturbed. This is because they feel no need to deny the Right Reasons Ability Condition. They can accept it as long as they hold a liberal view of what it takes to have the ability to act for the right reasons. In this subsection I will argue that in order to act for the right reasons, one must possess those reasons.

The easiest way to see why possession is necessary is by considering pairs of cases. Delusional Andy and Surprised Andy is one such pair.[13]

Delusional Andy

> Andy knows that his wife has always been an extremely loyal person. He also knows that he has no reason to think that she is cheating on him. Despite this

[13] Similar intuitions are pumped for similar reasons in Hyman (2006); Hornsby (2008); Gibbons (2001); Marcus (2012).

knowledge, he does believe that she is cheating on him. He thus moves out and files for divorce. In fact, his wife is cheating on him.

Surprised Andy

Andy knows that his wife has always been an extremely loyal person. However, much to his surprise, he learns that she is cheating on him—her best friend tells him, he finds some love letters, and he catches his wife with her lover. He thus moves out and files for divorce.

In both cases, Andy reasons from a belief that his wife is cheating on him to two actions—viz. moving out and filing for divorce. Indeed, we can suppose that delusional Andy's deliberation is phenomenologically indistinguishable from surprised Andy's. Moreover, in both cases Andy's belief is true. Finally, the fact that his wife is cheating on him is a weighty reason to perform those actions. However, it's very plausible that only in Surprised Andy does Andy move out and file for divorce *because* his wife is cheating on him. That is, it's very plausible that only Andy in Surprised Andy acts for the right reasons.

We have several ways of testing this intuition. First, we have Credit. If Credit is true and Andy acts for the right reasons in Delusional Andy, then Andy's token acts of moving out and filing for divorce should be creditworthy. But they intuitively aren't. Andy is delusional in Delusional Andy! Despite the fact that he performs the best action, he does not deserve credit for it. This is because he is just lucky that the act he actually performed turned out to be the best one.

Second, we have the Sensitivity Condition. If Andy acts for the right reasons in Delusional Andy, then he is sensitive to the fact that his wife is cheating on him is a reason to perform the actions he does perform. Intuitively Andy is not sensitive to that fact in the right way. He has no legitimate contact with that fact. He is just lucky that his irrational belief happens to be true. Because of this, it is hard to see how he is sensitive to that fact in the right way.

Third, we have the Explanation Condition. If Andy acts for the right reasons in Delusional Andy, then those reasons explain his moving out and filing for divorce. But this doesn't seem right, either. The fact that his wife is cheating on him doesn't seem to be the fact that explains his actions, given that he is untethered from that fact. Since he isn't sensitive to that fact, it doesn't seem like it can explain his actions.

We should come to the opposite conclusions about Andy in Surprised Andy. His actions in that case do seem creditworthy; they do seem sensitive to the relevant facts; and they do seem to be explained by the fact that his wife is cheating on him. The only relevant difference between the two Andys is that in Surprised Andy the relevant fact is within Andy's epistemic ken. Thus, it seems like in order to act for the right reasons, those reasons have to be within your ken. In other words, in order to act for the right reasons, you have to possess those reasons.

At this point we should consider an important objection to this defense of (2). The anchor of the objection is the obvious fact that even in Delusional Andy, the consideration that

Andy's wife is cheating on him plays an important role in Andy's deliberation and subsequent action. There is a sense in which the reason for which Andy moves out and files for divorce is that his wife is cheating on him. After all, that is the thought that ultimately motivates Delusional Andy to move out and file for divorce. Why not think that this is sufficient for meeting the Right Reasons Ability Condition? Isn't having the ability to be motivated by the relevant considerations all that is required?

I agree that there is a sense in which the reason for which Andy moves out and files for divorce is that his wife is cheating on him. We can make his actions *intelligible* by citing that consideration. We can, that is, understand why he did those things rather than some other things by appealing to the (content of) his belief that his wife is cheating on him. And we can do this even though he is delusional.

As it happens, most theories of acting for reasons are theories of intelligibility, in this sense.[14] These theories seek to understand what the two Andys have in common. Moreover, they all hold that delusional Andy's belief that his wife is cheating on him plays an important part in explaining why he acted as he did.[15] So it is not unmotivated to think that there is a sense in which the reason for which delusional Andy acts is that his wife is cheating on him.

The question now is whether acting for a consideration that happens to be the right reason is sufficient for acting for the right reasons. If it is, then delusional Andy does act for the right reason. And if this is right, then the objectivist can deny (2) while retaining (1).

Not surprisingly, I don't think acting for a consideration that happens to be the right reason is sufficient for acting for the right reasons. The problem is that intelligibility is a very low threshold. It is because of this that it is implausible that being able to make one's actions intelligible by citing a consideration that happens to be the right reason is sufficient for acting for the right reasons.

Perhaps the most straightforward way of seeing this is by considering cases where one's action can be made intelligible by citing a consideration that aligns with a normative reason, but one doesn't treat that consideration in the way that it ought to be treated. John's Home is one such case.

John's Home

> John stayed home from work sick. In fact, he is very sick. He badly needs medical attention. Becky is John's one time friend turned arch nemesis. She finds out he is in need of medical attention. Given that he is in a compromised physical state, she decides to go to his house to complete her plan of murdering him.

[14] To name just a few: Davidson (1980); Dancy (2000); Smith (1994); Setiya (2007). For some pushback on the thought that we can assimilate the two Andys, see Gibbons (2001); Hornsby (2008); Lord (2013); Lord & Sylvan (MS).

[15] Of course, the role it plays is different in different theories. I'll be assuming something like Dancy's view, which holds that the consideration itself is the motivating reason.

In John's Home, the fact that John badly needs medical attention is a strong reason to go to John's house. Moreover, Becky does go to John's house, and this action can be made intelligible by citing the fact that John badly needs medical attention. However, Becky clearly doesn't act for the right reason. She goes to John's house not to provide him with medical attention, but to murder him. Thus, acting for a consideration that aligns with the right reason is not sufficient for acting for the right reasons.

Granted, there is a disanalogy between Delusional Andy and John's Home. On the most natural interpretation of Delusional Andy, Andy treats the consideration that his wife is cheating on him as a strong reason to move out and file divorce. Thus, he treats that consideration in the right way, which is to say that he at least implicitly has the right view about the normative relevance of the fact that his wife is cheating on him. This isn't true in John's Home. Becky doesn't treat the relevant considerations in the 'right way.' She treats the fact that he is badly in need of medical attention as a reason to go to his house because she wants to kill him, not help him. This doesn't spoil the point, though. The point was just that having the ability to act in a way that enables one to make your act intelligible by citing a particular consideration is not sufficient for having the ability laid down by the Right Reasons Ability Condition.

That said, cases like John's Home suggest that perhaps we just need to add that one treats the consideration in the right way, whatever that turns out to be (we can rely on an intuitive understanding). I don't think this will do, either. The following variation on Mine Shaft makes this vivid.

Random Picking

> Everything that is true in Mine Shaft is true in this case. Billy knows that if the miners are in *A*, then it is best to sandbag *A* and she knows that if the miners are in *B*, it is best to sandbag *B*. Moreover, she is disposed to treat those considerations in the right way. She is disposed to sandbag *A* if she believes the miners are in *A* and disposed to sandbag *B* if the miners are in *B*. She has also recently taken a pill that will arbitrarily cause her to either believe the miners are in *A* or the miners are in *B*. The pill kicks in and she believes the miners are in *A*. She acts on this belief and sandbags *A*. As it happens, the miners are in *A*.

Clearly Billy doesn't act for the right reason in this case. However, her act can be made intelligible by citing the consideration that the miners are in shaft *A*. Moreover, she treats this consideration in the right way. Indeed, she isn't that much different from Delusional Andy. They both reason—in the right way—from an paradigmatically irrational belief to an action. It seems clear upon reflection that they lack the ability to act for the right reasons.

The important feature of cases like Delusional Andy and Random Picking is that the characters' beliefs lack a certain epistemic pedigree. Thus far we haven't confirmed that this the same pedigree involved with possession. We have ruled out, I take it, that merely believing some proposition and acting on it in the right way is sufficient for acting for the right rea-

sons.[16] What seems to be required is that agents stand in some positive epistemic relation to the reason.[17]

To argue for any specific view about which positive epistemic relation is involved would be to break my neutrality about what the possession relation involves. I don't think this is necessary here. I am content to leave it here: It is very plausible that possession requires a positive epistemic pedigree and it is very plausible that acting for the right reason requires a positive epistemic pedigree. It would be quite odd if it wasn't the same epistemic pedigree. I think this is enough to abductively infer that in order to meet the Right Reasons Ability Condition, one must possess the right reasons.[18] In other words, (2) is true.

From (1) and (2) we can infer that if you ought to ϕ, then you possess the right reasons to ϕ. Since possession requires that the reason be within your perspective, it follows that objectivism is false.

This argument doesn't immediately establish the view that the reasons you possess determine what you're obligated to do. It just establishes that possession is necessary for a reason to obligate. It doesn't establish that possessing a set of reasons that conclusively support ϕ-ing is sufficient for those reasons to obligate you to ϕ. While it doesn't establish this, it is telling, I think. The most natural view to take once you're on board up to this point is that the reasons you possess determine your obligations. I will rest content with establishing the necessary condition and hence showing that objectivism is false.

4 New Information Problems

My argument for perspectivalism notwithstanding, there are some powerful arguments in the literature for objectivism. In the last half of the paper I will provide new perspectival replies to these arguments. I think that my particular perspectival view has the resources to adequately respond.

4.1 Two New Information Problems

The Past Obligations Problem

Suppose that right before the moment of truth, Billy figures out a way to determine where the miners are. She takes advantage of this procedure and thus finds out that the miners are in shaft A. Question: What should she think she was obligated to do before she found out this information? It is quite plausible that she should think that her obligations didn't change.

[16] I think this result can be parlayed into an argument against certain views of possession. See (Lord, 2013, ch. 2) for more on this.

[17] This isn't all that is required. You also have to treat it in the right way.

[18] I should stress that I don't think the story ends here. I go much further in (Lord, 2013, ch. 2). But in order for the story to continue I have to take sides on the nature of possession. I would like the results to be as general as possible and thus don't want to take sides here about possession.

Rather, she discovered what her obligations were by discovering where the miners were.[19] As (Ross, 1930, pg. 32) puts it, 'Many people would be inclined to say that the right act for me is ... that which on all the evidence available to me I should think to be my duty. But suppose that from the state of partial knowledge in which I think act A to be my duty, I could pass to a state of perfect knowledge in which I saw act B to be my duty, should I not say 'act B was the right act for me to do'?'

This thought is in tension with perspectivalism. This is because the perspectivalist is committed to thinking that before she got the new information, Billy's perspective best supported doing nothing. Moreover, when she gained the new information, her perspective best supported blocking *A*. Thus, it seems that Billy must be mistaken if she thinks that her obligation at the earlier time was the same as her obligation at the later time.

It seems that we can confirm that my particular perspectival view makes these predictions. Before finding out where the miners are, the reasons she possesses conclusively support doing nothing. After she finds out, this is no longer true. After she finds out the reasons she possesses decisively support blocking shaft *A*. So if Billy thinks she's discovered what her obligation was all along, she is mistaken, according to my perspectival view.

The Advice Problem

When we seek advice about what to do (and when we give advice about what to do), we don't seek advice about what our perspectives best support. We want to know what's best, not what's best given our perspective. To see this, suppose that Billy's method of figuring out where the miners are is to ask you what she ought to do. She has found out, suppose, that you know where the miners are. It would be a mistake for you to tell her she ought to do nothing. You should tell her to block shaft *A*. As (Thomson, 1986, pg. 179) bitingly puts a similar idea, 'On the rare occasions someone conceives of asking my advice on moral matter, I do not take my field work to be limited to a study of what he believes is the case: I take it incumbent upon me to find out what *is* the case.' Similarly, we might think, an advisor's field work (whether it's a moral matter or not) is not limited to what the evidence or knowledge of the advisee suggests is best; it's what is best.

Again, this seems to be in tension with perspectivalism. If perspectivalism is correct, then Billy ought to do nothing. So if she asks you what she ought to do, you should tell her she ought to do nothing. Since you should in fact tell her she ought to block *A*, it is very plausible that she ought to block *A*. Thus, perspectivalism must be false.

Once again, it seems as if my perspectival view makes the bad predictions. Before Billy gets the new information, the reasons she possesses conclusively support doing nothing. So this seems to be the answer to give when she asks what she ought to do. But this is not what you should say. You should tell her she ought to block shaft *A*.

[19] This is often taken as data, but not always. Some have argued that these type of hindsight judgments are incorrect. See, e.g., Bjornsson & Finlay (2010) and especially Dowell (FC). I have the intuition that the hindsight judgments are true, even though upon reflection I find them deeply puzzling.

4.2 A Diagnosis

I grant that these arguments have great appeal. What I want to know right now is why they have such appeal. For it is quite puzzling, to me at least, that one's intuitions about sophisticated ignorance cases can be tossed and turned so easily. It is very plausible, even upon reflection, to think that Billy ought to do nothing when she is ignorant. However, it also seems plausible that if she is relieved of her ignorance, her judgments about what she was obligated to do while ignorant should match her judgments about what she ought to do with more information. And it is especially plausible that when giving advice one should not always focus on the epistemic situation of the agent one is giving advice to—even when one is giving advice to a character in a sophisticated ignorance case. What gives?

What we want first is an explanation for why it is our hindsight judgments and advice giving practices behave this way. I think they do because deliberation aims at what's best or what's supported by all the reasons. This is why Billy thinks that her past obligation is the same as her current one—because her current one is getting at what she was aiming at all along. Moreover, it is plausible that advice is parasitic on our deliberative aims. That is, correct advice is guided by the aims of deliberation. This is why we seek to inform the advisee what's best or what's supported by the balance of all the reasons when we give advice.

The million dollar question is whether my perspectival view is compatible with the claim that deliberation aims at what's best. I think that it is. Indeed, I think that it is also compatible with thinking that Billy's hindsight judgment and your advice about what Billy ought to do are true.

To stress, I think there are two burdens here. The first is to explain what is going on with Billy's hindsight judgment and with the advisor's assertion. Are they true? Is this compatible with perspectivalism? The second burden is to show that perspectivalism is compatible with the intuition that motivates our judgments—viz., that deliberation aims at what's best. I'm going to take these in reverse order. I will first argue that my perspectivalist view is compatible with the claim that deliberation aims at what's best. I will then explain why I think the hindsight judgment and the advice are true and how this relates to perspectivalism.

4.3 Abilities, Obligation, and 'Ought'

My solution to the new information problems has two parts. The first part is about the *metaphysics* of obligation. The second part is about the *semantics* of 'ought' in English. I think that the two come apart in predictable ways given the aim of deliberation. When it comes to the metaphysics, my perspectival view is compatible with thinking that the aim of deliberation is to do what's best. But our 'ought' thought and talk tends to track the aim of deliberation, so to speak, and in contexts where one party has more information this will lead us to make judgments about our obligations that come apart from our deliberative obligations. This is why the truth-values of the relevant sentences in English relevant to the evaluation of the hindsight judgment and advice come out true. Let's start with the metaphysics.

New Information as a Problem for the Metaphysics

The first point to make is that it is not at all clear that the new information problems speak decisively in favor of objectivism. This is because there are cases that provide the same lesson even though the new information provided doesn't put us in a position to know what is best or supported by all the reasons. Consider a version of Mine Shaft where Billy starts out being more ignorant than in the original case (new details in bold).

More Ignorant Mine Shaft

> A group of 10 miners are trapped in a mine. They are either trapped in shaft *A* or in shaft *B*. It is not known which shaft they are in. Flood waters are approaching the shafts. Billy has the choice to sandbag shaft *A*, sandbag shaft *B*, or not sandbag either. She knows that if she sandbags *A* and the miners are in *A*, all the miners will survive. **She has strong but misleading evidence, however, that if she sandbags *A* and they aren't in *A*, nothing bad will happen and vice versa for getting it wrong about *B*.** Finally, she knows that if she does nothing, then 9 of the 10 will survive.

Given what Billy knows, doing nothing is definitely not what she ought to do. She ought to block *A* or *B*. Now imagine that you know that if she blocks *A* and the miners are in *B*, then all the miners will die and you know that if she blocks *B* and the miners are in *A*, all will die. If she asks you what the thing to do is, you should tell her to block neither shaft, even though what she ought to do given her perspective is block *A* or *B*.

We can say similar things about Billy's thoughts about past obligations. Suppose you tell her the new information and it becomes true that from her perspective she ought to block neither shaft. It would be natural for Billy to think that she's discovered what she ought to have done all along. She will thus judge that her earlier thought about what she ought to do was false.[20]

The rub, of course, is that in this case both you and Billy *know* that blocking neither shaft is not best. So it seems that the new information problems don't show that the nature of deliberation and advice entail or even support objectivism. It seems as if the lessons can be learned by focusing on different sophisticated ignorance cases. Back to this in a moment.

The second point to make is that it is very plausible that ability conditions generally are compatible with the thought that deliberation aims at what's best. This is obvious when it comes to the Physiological Ability Condition. Deliberation can aim at what's best even though our obligations are *constrained* by our physiological abilities. That is, deliberation can aim at what's best even though we aren't always obligated to bring about the best state of affairs because sometimes we don't have the physiological ability to bring about the best state of affairs. This much is obvious.

[20] I imagine many who have a strong intuition in the original hindsight don't have as strong an inutition in this case. I think this is some reason to doubt the veridicality of our intuition about the original case. Again, though, I will grant the data for the sake of argument.

I think the same is true of the Right Reasons Ability Condition. Deliberation can aim at what's best even though our obligations are constrained by some of our agential abilities. That is, deliberation can aim at what's best even though we aren't always obligated to bring about the best state of affairs because sometimes we don't have the agential abilities needed to bring about the best state of affairs in a way deserving of credit.

So far we've seen that it is intelligible to think that my perspectival view is compatible with thinking deliberation aims at what's best, but we haven't been told explicitly why we should think this is true. I think that cases like More Ignorant Mine Shaft provide some strong evidence that we implicitly recognize the relevant constraints. I see no reason to think that in More Ignorant Mine Shaft Billy doesn't seek what's best in her deliberation. Nor is there any reason to think that you, her advisor, are eschewing the aim of having Billy do what's best. However, you recognize that pursuit of that aim is constrained by the information within your perspective.

It's helpful here to compare practical deliberation with *epistemic deliberation* and its aims. Plausibly, epistemic deliberation—deliberation about what to believe—aims at the truth. Given this, you'd expect there to be a new information argument for the conclusion that one is always deliberatively obligated to believe the truth. At the very least, epistemic advisors try to advise their advisees to believe truths. And in cases where the advisee's information suggests $\neg p$ and the advisor's better information suggests p, the good advisor should tell the agent they ought to believe p. Does this show that we're always deliberatively obligated to believe the truth?

No, this argument is obviously bad. It is incredibly plausible that we are sometimes deliberatively obligated to refrain from believing the truth. There are two relevant cases. In the first, we are deliberatively obligated to believe something that is false. Sometimes the evidence available is misleading and strongly supports believing p even though $\neg p$. In these cases, it's plausible that we are deliberatively obligated to believe p.

Even if you think that we are never deliberatively obligated to believe a falsehood,[21] it is still overwhelmingly plausible that we are sometimes not deliberatively obligated to believe the truth. This is because we are sometimes deliberatively obligated to *withhold belief*. To take the easiest case, when we lack evidence for both p and $\neg p$, we deliberatively ought to withhold belief about p. Consider an example. There is a grassy field in one corner of Central Park. The number of blades of grass in this field is either odd or even. Consider with me the question of whether the number is odd or even. If you're like me, you have *no evidence* either way. Because of this, it is very plausible to think that we ought to withhold belief on this question. If the new information argument for objectivism worked, it seems like what we in fact ought to do is either believe it is odd or believe it is even.[22]

[21] If you are, for example, Littlejohn (2012).

[22] Notice that cases where we ought to withhold are *just like* sophisticated ignorance cases. That is, they are cases where we are in a position to know that the option that we ought to take is second best. Given how plausible it is that this is the right answer in the epistemic case, we should be more confident that the sophisticated ignorance cases are indeed counterexamples to objectivism.

It is worth emphasizing the implausibility of this result. If epistemic deliberation aims at the truth and the new information argument is sound, then we are forced to think that we are always obligated to believe the truth. One result of this is that we are always obligated *not* to withhold since withholding is incompatible with believing. This is extremely implausible.

Most importantly for our purposes, it doesn't seem like obligations to withhold threaten the thought that epistemic deliberation aims at the truth. What we're after is the truth even though sometimes the only permitted option is to withhold. Intuitively, this is because we are only allowed to rely on certain information in deliberation and, alas, sometimes that information doesn't adequately support either p or $\neg p$. In other words, it seems very plausible that our epistemic obligations are constrained by our perspectives. This does not seem to threaten the thought that epistemic deliberation aims at the truth.

To be clear, there are two lessons to be learned from the epistemic. The first is that perspectivalism seems incredibly plausible in the case of epistemic obligation. This in itself lends support to perspectivalism about practical obligation, for without some story about what's different between the epistemic and the practical, what goes for the epistemic should go for the practical.

This first lesson, while important, isn't the most important lesson for my purposes here. The most important lesson for my purposes here is that in the epistemic case it is very plausible that we aren't always obligated to believe the truth even though epistemic deliberation aims at the truth. There seem to be constraints on our epistemic obligations. This is the structure I am suggesting practical deliberation and obligation have. The epistemic case provides a nice model of how I think the practical works. Appreciating how natural the structure is in the epistemic case is the main lesson I want to glean from the epistemic.

Thus, I don't think that the new information problems present much of a problem for the metaphysics of my perspectival view of obligation. This is because my view can account for the claim that motivates the new information problems, which is the claim that deliberation aims at what's best. Deliberation aims at what's best even though our obligations are constrained in various ways by our abilities.

New Information and the Semantics of 'Ought'

While the metaphysical problem is, I take it, the biggest problem posed by new information, it isn't the only problem. This is because so far we haven't accounted for the data in the new information cases.[23] Namely, we haven't accounted for the fact (if it is a fact) that Billy's judgment about her past obligation—i.e. her judgment that her past obligation is the same as her obligation after gaining more information—is true. We also haven't explained how it is that you—Billy's advisor—say something true when you tell her (e.g.) that she ought to block shaft A.

[23] Again, it's important to stress that I'm not completely sold on the hindsight cases. I am merely granting it is data for the sake of argument.

It certainly seems like these bullets must just be bitten. That is, it seems like it's not possible for the perspectivalist to account for the above data. I think these seemings are misleading. In order to see why we need to understand a bit about how 'ought' works in English.

On the canonical view of 'ought' in linguistics, 'ought' operates as a quantifier over possible worlds.[24] But it doesn't (always) quantify over all possible worlds. Rather, it quantifies over a restricted set of possible worlds. Which set it quantifies over is determined by context. One way in which context often restricts the domain is by restricting the amount of information that can be taken as true. The most natural way this happens is by limiting the domain to the worlds compatible with some salient body of information. Sometimes this is just the knowledge of the speaker and sometimes it's the knowledge that a group of contextually salient speakers has.[25] When the domain gets restricted in this way, we can have thoughts and talk about what one ought to do given some limited body of information.[26] We often take advantage of this nice feature of the word 'ought.'

Given the role that context plays in the semantics of 'ought,' 'ought' thought and talk is rather flexible. We can think and talk about what ought to be done given X for a very large amount of Xs. This means that Billy can think about her past obligations in light of her new information. Moreover, it means that advisors can think about the obligations of advisees from the perspective of their information. I think these are the contents of Billy's thought and of the advisor's thought (and assertion). And surely those contents are true.

Doesn't that show that perspectivalism is false? In a word: No. As we've already seen, there are lots of true 'ought' claims in this case. It's true relative to some bodies of information that Billy ought to do nothing, it's true relative to some bodies of information that she ought block shaft A and it's true relative to other bodies of information that she ought block A or block B. Those truths don't necessarily establish anything about what she deliberatively ought to do.

'Fair enough,' one might respond, 'but this leaves out the important fact that Billy and the advisor are having those true thoughts in a deliberative context and, moreover, the content of Billy and the advisor's thoughts seems to be the answer to the central deliberative question. This provides very strong evidence that Billy and the advisor really are getting at the deliber-

[24] It is the canonical view mostly because of the pioneering work of Angelika Kratzer. See Kratzer (2012). Recently there has been much debate about the exact role these relativizations play in the semantics. Contextualists like Kratzer think that the relativization plays a role in determining the content of the propositions expressed, whereas truth-relativists like Kolodny & MacFarlane (2010) hold that the content is contextually invariant but that the truth value is relativized to contexts of assessment in another way. This debate is orthogonal to our discussion here. I can make my main point no matter who is right about how the relativization works in the semantics. I will assume contextualism, though, given that it is the canonical account.

[25] There is a second way that context can play a role in the semantics. Namely, by fixing which standards will be germane for the evaluation of the options. For example, sometimes the rules of etiquette will be selected, sometimes the law, sometimes morality etc. We can just ignore this here and assume that the standards chosen are the standards that evaluate the deliberative 'ought.'

[26] Sometimes this relativization to information will be explicitly contained in what's said, for example when we say things like 'Given what Billy knows, she ought to ϕ.' Most often, though, we just make bare 'ought' claims and context determines the relativization.

ative ought.'

This is a powerful response. However, I think it can be resisted. Those who want to resist it have at least two burdens. First, they have to explain why it is that Billy and the advisor's thoughts don't track Billy's deliberative obligations in these particular cases. But this isn't enough to be fully satisfying. For once we have this explanation, we'll want to know if it generalizes (or overgeneralizes). That is, we'll want to know whether the explanation, if correct, shows that we never or rarely track deliberative obligations. This would be bad. It's a very serious problem with a theory if it's committed to holding that our 'ought' thought and talk very rarely tracks our deliberative obligations. So this commitment should be resisted. If it is, then—and this is the second burden—one needs to explain why it is that in these cases we don't track our deliberative obligations but in most cases we do.

I think both burdens can be met. Let's start with the first: Why is it that Billy's and the advisor's thoughts don't track Billy's deliberative obligations? It is because, I think, Billy and her advisor are more concerned with what's best rather than what Billy's deliberative obligations are in a more ignorant state. This is not surprising given that the aim of deliberation is doing what's best. This is what we're trying to get at in deliberation. Given that, it is no surprise that the 'ought' judgments we are disposed to make will always be relativized to the best information available. What our deliberative obligations are given worse information is of no interest to us given our aims.[27]

This explanation clearly doesn't overgeneralize and thus the second burden can also be met. If my metaphysical story above is correct, then deliberation aims at what's best even though our obligations are constrained by some of our abilities. One of these constraints is tied to how much information one has. Given this, what's best *now* in light of the information *currently* had is of great interest to a deliberator (and the advisers of the deliberator). But true thoughts about this will track one's deliberative obligations. So the explanation of why Billy's thought about her past obligation and the advisor's thought don't track Billy's deliberative obligations doesn't overgeneralize. Since our obligations are constrained by the information we have and the information we have will be all deliberation can go on, most of our 'ought' thoughts will track our deliberative obligations. It is only when we get differences in how much information is possessed between agents (or time-slices of agents) that we get the two coming apart.

To recap: My response to the new information problems is two fold. First, I think that new information does not put pressure on my view of the metaphysics of obligation. This is because the motivating idea behind the arguments—that deliberation aims at what's best—is compatible with my perspectival view. Deliberation can aim at what's best even if obligation is constrained by our abilities.

The second part of my answer has to do with the semantics of 'ought' in English. Given the flexibility of 'ought,' we can have all kinds of true thoughts about what we ought to do. This means that our hindsight judgments can be about what we ought to have done given the

[27] Cf. Bjornsson & Finlay (2010).

information we have now and our advisors' thoughts can be about what we are obligated to do given their information. I think that these are the contents of our thoughts in hindsight cases and the content of our advisors' thoughts. And those contents are true. Moreover, there is a plausible story to be told about why our thinking about what ought to have been done in hindsight and our advisors' thinking about what we ought to do can come apart from thinking about our deliberative obligations.

5 Conclusion

This paper had two main ambitions. The first was to provide an argument for perspectivalism. The key idea of that argument was that in order for a reason to obligate, it has to be potentially action guiding in a certain way—it has to be possible for that to be the reason for which we act. I argued that a reason can be potentially action guiding in this way only if we possess that reason. Thus, perspectivalism is true and objectivism is false.

The second ambition was to respond to what I take to be the strongest argument against perspectivalism in favor of objectivism. I argued that the motivating thought behind that argument is compatible with my perspectival view. Moreover, I provided explanations of the key data that are both compatible and friendly to my view.

References

Bjornsson, G. & Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121(1), 7--36. [cited on pg. 16, 22]

Dancy, J. (2000). *Practical Reality*. Oxford University Press. [cited on pg. 13]

Davidson, D. (1980). *Essays on Actions and Events*. Oxford University Press. [cited on pg. 13]

Dowell, J. (FC). Flexible contextualism about deontic modals. *Inquiry*. [cited on pg. 16]

Gibbons, J. (2001). Knowledge in action. *Philosophy and Phenomenological Research*, 62(3), 579--600. [cited on pg. 11, 13]

Graham, P. A. (2010). In defense of objectivism about moral obligation. *Ethics*, 121(1), 88--115. [cited on pg. 1, 2, 6]

Hornsby, J. (2008). A disjunctive conception of acting for reasons. In A. Haddock & F. Macpherson (Eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford University Press. [cited on pg. 11, 13]

Hyman, J. (2006). Knowledge and evidence. *Mind*, 115(460), 891--916. [cited on pg. 11]

Jackson, F. (1991). Deontic-theoretic consequentialism and the nearest-dearest objection. *Ethics*, 101(3), 461--482. [cited on pg. 5, 6]

Kiesewetter, B. (2011). "Ought" and the perspective of the agent. *Journal of Ethics & Social Philosophy*. [cited on pg. 6]

Kolodny, N. & MacFarlane, J. (2010). Ifs and oughts. *The Journal of Philosophy*, 107(3), 115--143. [cited on pg. 6, 21]

Kratzer, A. (2012). *Modals and Conditionals: New and Revised Perspectives*. Oxford University Press. [cited on pg. 21]

Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge University Press. [cited on pg. 19]

Lord, E. (2013). *The Importance of Being Rational*. PhD thesis, Princeton University. [cited on pg. 4, 8, 13, 15]

Lord, E. & Sylvan, K. (MS). Prime time (for the basing relation). Manuscript, Princeton University and Rutgers University. [cited on pg. 8, 13]

Marcus, E. (2012). *Rational Causation*. Harvard University Press. [cited on pg. 11]

Moore, G. E. (1912). *Ethics*. Oxford University Press. [cited on pg. 1, 2]

Parfit, D. (2011). *On What Matters*. Oxford University Press. [cited on pg. 5]

Regan, D. (1980). *Utilitarianism and Cooperation*. Oxford University Press. [cited on pg. 5]

Ross, J. (2006). *Acceptance and Practical Reason*. PhD thesis, Rutgers University. [cited on pg. 5]

Ross, W. (2002 (1930)). *The Right and the Good*. Oxford University Press. [cited on pg. 2, 16]

Schroeder, M. (2009). Means-end coherence, stringency, and subjective reasons. *Philosophical Studies*, 143(2), 223--248. [cited on pg. 6]

Setiya, K. (2007). *Reasons Without Rationalism*. Princeton University Press. [cited on pg. 13]

Smith, M. (1994). *The Moral Problem*. Blackwell Publishing. [cited on pg. 13]

Thomson, J. J. (1986). Imposing risks. In *Rights, Restitution, and Risk*. Harvard University Press. [cited on pg. 1, 2, 16]

Wallace, R. J. (2012). Duties of love. *Aristotelian Society Supplementary Volume*, 86(1), 175--198. [cited on pg. 1]

Zimmerman, M. J. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press. [cited on pg. 6]